Dynamic Asset Allocation Approaches for Counter-Piracy Operations

Woosun An, Diego Fernando Martinez Ayala, David Sidoti, Manisha Mishra, Xu Han, and *Krishna R. Pattipati
Dept. Electrical and Computer Engineering
University of Connecticut
Connecticut, United States
*krishna@engr.uconn.edu

[1]Eva D. Regnier, [1]David L. Kleinman, and [2]James A. Hansen
[1]Naval Postgraduate School
California, United States
eregnier@nps.edu and dlkleinm@nps.edu
[2]Naval Research Laboratory
California, United States
james.hansen@nrlmry.navy.mil

Abstract — **Piracy on the high seas is a problem of world-wide concern. In response to this threat, the US Navy has developed a visualization tool known as the Pirate Attack Risk Surface (PARS) that integrates intelligence data, commercial shipping routes, and meteorological and oceanographic (METOC) information to predict regions where pirates may be present and where they may strike next. This paper proposes an algorithmic augmentation or add-on to PARS that allocates interdiction and surveillance assets so as to minimize the likelihood of a successful pirate attack over a fixed planning horizon. This augmentation, viewed as a tool for human planners, can be mapped closely to the decision support layer of the Battlespace on Demand (BonD) framework [32]. Our solution approach decomposes this NP-hard optimization problem into two sequential phases. In Phase I, we solve the problem of allocating only the interdiction assets, such that regions with high cumulative probability of attack over the planning horizon are maximally covered. In Phase II, we solve the surveillance problem, where the area not covered by interdiction assets is partitioned into non-overlapping search regions (e.g., rectangular boxes) and assigned to a set of surveillance assets to maximize the cumulative detection probability over the planning horizon. In order to overcome the curse of dimensionality associated with Dynamic Programming (DP), we propose a Gauss-Seidel algorithm coupled with a rollout strategy for the interdiction problem. For the surveillance problem, we propose a partitioning algorithm coupled with an asymmetric assignment algorithm for allocating assets to the partitioned regions. Once the surveillance assets are assigned to search regions, the search path for each asset is determined based on a specific search strategy. The proposed algorithms are illustrated using a hypothetical scenario for conducting counter-piracy operations in a given Area of Responsibility (AOR).**

Keywords-component: *Resource management problem, Search problem, Partitioning algorithm, Approximate dynamic programming, Allocation problem, Rollout, Gauss-Seidel iteration*

I. INTRODUCTION

A. Motivation

The United States Navy faces a number of asymmetric threats (e.g. terrorists, pirates, drug smugglers) characterized by multiple illicit agents whose locations are generally unknown and whose behavior is generally unpredictable. A response to these threats requires: 1) integration of intelligence and effective surveillance to detect and identify threats in order to gain situational awareness, followed by 2) effective allocation of resources for interdicting the potential threats. This is the two-pronged problem that we address in this paper.

Recently, piracy in the Somali Basin and Gulf of Aden (GOA) has become a major international problem. According to International Maritime Bureau's piracy reporting center, there has been a significant increase in total number of pirate attacks in recent years (239 attacks in 2006 compared to 439 attacks in 2011) [29][30]. An increase in piracy activities has spurred the US Navy to develop a software model that integrates classified intelligence data, commercial shipping routes, and environmental information (e.g., wind speeds and direction, wave heights, and ocean currents) to predict where the pirates may be present and where they may strike next [1]. The model outputs consist of a set of color-coded maps designated the Pirate Attack Risk Surface, herein referred to as PARS [12][13]. For each forecast period, the ocean surface is partitioned into geographical "cells," and the PARS map predicts the probability of pirate attack for each cell taking into account intelligence, known pirate behavior, commercial shipping patterns, and weather patterns that may affect the pirates' ability to operate on small skiffs with the intent of attacking commercial ships. The PARS is updated every 12 hours, or when new intelligence comes in. Multinational counter-piracy forces operating in the region seek to deter and interdict pirate attacks, and should nominally have access to the PARS information. Indeed, the U.S. Naval Forces Central Command refers to the PARS product daily. Our goal is to augment PARS with asset allocation tools to assist human decision makers (DMs) involved in counter-piracy planning.

In general, the counter-piracy mission involves surveillance *and* interdiction operations. As shown in Fig. 1, the DMs choose from a set of available interdiction assets and provide commands, or a Course of Action (COA), for positioning these assets over a near-time planning horizon. The DMs allocate the interdiction assets such that regions with high probability of attack are maximally covered to neutralize and mitigate pirate attacks, while at the same time identifying the pirates [14]. Here, the surveillance assets (e.g., P3s) are also controlled by DMs; these assets are generally used for large ocean

Report Documentation Page

1. REPORT DATE **JUL 2012**	2. REPORT TYPE	3. DATES COVERED **00-00-2012 to 00-00-2012**

4. TITLE AND SUBTITLE **Dynamic Asset Allocation Approaches for Counter-Piracy Operations**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Naval Postgraduate School,Monterey,CA,93943**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**

13. SUPPLEMENTARY NOTES **Presented at the 15th International Conference on Information Fusion held in Sinapore on 9-12 July 2012. Sponsored in part by Office of Naval Research and Office of Naval Research Global.**

14. ABSTRACT
Piracy on the high seas is a problem of world-wide concern. In response to this threat, the US Navy has developed a visualization tool known as the Pirate Attack Risk Surface (PARS) that integrates intelligence data, commercial shipping routes, and meteorological and oceanographic (METOC) information to predict regions where pirates may be present and where they may strike next. This paper proposes an algorithmic augmentation or add-on to PARS that allocates interdiction and surveillance assets so as to minimize the likelihood of a successful pirate attack over a fixed planning horizon. This augmentation viewed as a tool for human planners, can be mapped closely to the decision support layer of the Battlespace on Demand (BonD) framework [32]. Our solution approach decomposes this NPhard optimization problem into two sequential phases. In Phase I, we solve the problem of allocating only the interdiction assets such that regions with high cumulative probability of attack over the planning horizon are maximally covered. In Phase II, we solve the surveillance problem, where the area not covered by interdiction assets is partitioned into non-overlapping search regions (e.g., rectangular boxes) and assigned to a set of surveillance assets to maximize the cumulative detection probability over the planning horizon. In order to overcome the curse of dimensionality associated with Dynamic Programming (DP), we propose a Gauss-Seidel algorithm coupled with a rollout strategy for the interdiction problem. For the surveillance problem, we propose a partitioning algorithm coupled with an asymmetric assignment algorithm for allocating assets to the partitioned regions. Once the surveillance assets are assigned to search regions, the search path for each asset is determined based on a specific search strategy. The proposed algorithms are illustrated using a hypothetical scenario for conducting counterpiracy operations in a given Area of Responsibility (AOR).

15. SUBJECT TERMS

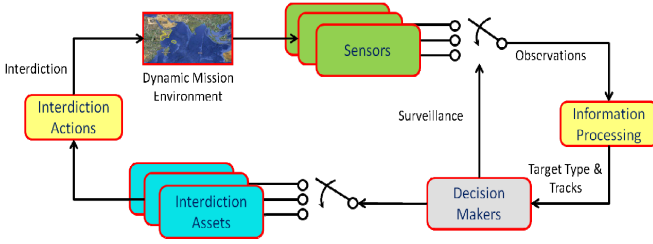| 16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**8**	

Figure 1: Counter-piracy problem as a stochastic control problem

surveillance and are assigned to predefined search regions, typically in the shape of "rectangular boxes". Observations from these assets are processed to characterize the target types and their trajectories. This information is relayed back to the DMs so that they can adjust their surveillance and interdiction plans for the next time interval. This planning process is then repeated, typically on a 12-24 hour cycle.

The combined surveillance and interdiction problem may be viewed as a moving horizon stochastic control problem as shown in Fig. 1. The key problem is to efficiently allocate a set of heterogeneous sensing and interdiction assets to minimize the likelihood of a successful pirate attack, subject to mission constraints such as the weather, asset availability, asset capabilities (e.g., range, speed), and asset assignment (e.g., many sensors may need to be coordinated to obtain a better picture of the situation).

B. Previous Work

Searching for targets (surveillance in a bounded region) is one of the most important applications of dynamic resource management. In general, the search problem can be categorized based on the number of targets and the assets involved. These problems come under the rubric of Partially Observed Markov (Semi-Markov) Decision Processes (POMDP (POSMDP)).

For a single searcher-single target problem, Eagle [18] considered a model wherein the searcher's movements were constrained only to the cells adjacent to the currently assigned cell (e.g., move, up, down, left, right or stay put). Eagle [18]-[19] found an optimal solution to a small search problem using a dynamic programming (DP) technique; however, the DP technique is infeasible for large search regions (>20-30 cells). Martins [7] introduced a branch-and-bound procedure, using the expected number of detections (ED) as a metric to be maximized; ED is an upper bound on the probability of detection (PD). In order to solve this problem, Martins created a network flow graph, wherein the arcs linking the cells corresponded to rewards and the longest path in this network corresponded to the ED bound. Lau [8] improved the bounds by tightening the ED bound with almost no added computation, viz., the so-called DMEAN bound. We derived a generalized mean (GMEAN) bound, which allows one to derive even tighter bounds than the DMEAN bound [24]. The search problems considered above assume that the false alarm probability is equal to zero. However, these problems are still difficult to solve as they are known to be NP-hard [19].

For a single target with multiple searchers, Santos [20] developed heuristic and optimal approaches to solve the search problem assuming that the searchers could move only to the neighboring cells at each time epoch. In our previous work [9],

we formulated the ASW asset allocation and search path problem using a hidden Markov modeling (HMM) framework. A searcher can observe multiple cells and the searchers may probabilistically interfere with each other if they observe the same cells. In order to solve this NP-hard optimization problem, we used a greedy approach, based on the evolutionary algorithm, coupled with the auction algorithm and Voronoi tessellation approach for partitioning the search region [9],[24].

For multiple target search problems, it is computationally intractable to even update the belief map. Nair [11] considered a search problem with multiple searchers and an unknown, but fixed, number of stationary targets in a given region and presented a computationally tractable method to compute the belief map update using the theory of symmetric polynomials. Royset and Sato [10] formulated a multiple target search problem as a convex mixed-integer nonlinear program and solved it using two methods; a cutting-plane method, and a method based on linearization of the cost function. They assumed that the number of targets and distribution of target paths are known. In contrast, we assume that the PARS maps, which take into account intelligence, known pirate behavior, commercial shipping patterns, and weather patterns, encode the information state for asset allocation.

The counter-piracy problem, as a specific application of the search problem, has attracted much interest due to an increase in the number of pirate activities in recent years. Marsh [26] provided a game theoretic model, where one interdiction asset and one surveillance asset are utilized for a counter-piracy mission. Due to the two-person zero sum game structure, the model is limited to the case where pirates have complete knowledge, while the interdiction-surveillance asset combination has some predisposed beliefs.

Kress *et*. *al* [27] developed a stochastic dynamic programming model for the combined search and interdiction problem where a *single* surveillance and *single* interdiction asset is engaged in searching for, identifying and interdicting hostile vessels within a given time frame. Due to computational intractability, they proposed a greedy heuristic approach for solving this counter-piracy problem.

In this paper, we consider the problem of allocating *multiple* assets for surveillance and interdiction operations, to detect and interdict *multiple* targets within a vast area of responsibility (AOR). We decompose our problem into two sequential phases by exploiting the fact that interdiction assets (ships that may have helicopters on board) are substantially slower than the surveillance assets (e.g., P-3 aircraft). In Phase I, we solve the problem of allocating the interdiction assets with different capabilities (e.g., the reachable cells and interdiction range per unit time interval). In order to overcome the curse of dimensionality associated with DP, we propose to combine the Gauss-Seidel approach (method of successive displacements) with rollout concepts.

The primary objective is to locate assets to deter or interdict attacks; it is assumed that if a vessel is being harassed by pirates it will call for help so that a nearby asset can deter the attack. In this model, it is assumed that it is not necessary to detect pirates' precise location within the region covered by interdiction assets. Surveillance is valuable for updating

information about pirates' locations for future assignment of interdiction assets. In Phase II, the area not covered by interdiction assets is partitioned into non-overlapping search regions (rectangular boxes). However, this problem is also NP-hard and we propose a partitioning algorithm, where each region, starting with a single cell, is grown independently subject to the region's shape constraints, and couple it with an asymmetric assignment algorithm for allocating surveillance assets to the partitioned regions. Once the surveillance assets are assigned, the search path for each asset is computed by tactical units (e.g., individual P-3) assigned to the search regions. The capabilities of the proposed algorithms are illustrated using a hypothetical scenario to detect and interdict pirate activities in an AOR.

C. Organization of the Paper

The paper is organized as follows. In section II, we formulate the interdiction problem and propose a Gauss-Seidel method for its solution. In section III, we formulate the surveillance problem and propose a combined partitioning and asymmetric assignment algorithm. In section IV, the capabilities of the proposed algorithms are illustrated using a hypothetical scenario. Section V concludes with our summary.

II. INTERDICTION PROBLEM

A. Problem Formulation

We define a time period k of length Δ, where Δ (e. g., 12 or 24hrs) is the time interval between command updates to the available assets. Assume that time period $k = 0$ is the current time period, $k = 1$ is the next time period, and so on. DMs plan at current time period ($k = 0$) for where assets are to be positioned for the next K periods, $k=1, 2,…, K$.

The set of interdiction assets that are assignable during period $k > 0$ is denoted as I_k, $k=1, 2 …K$. The AOR, G, is partitioned into cells denoted by $g \in G$. PARS, updated every time period, provides the probability of pirate attack in cell g during time period k, denoted by $PA(g, k)$. The cell location of asset i during time period k is denoted as $x_i(k)$. Decisions are made today as regards to the future positioning of assignable assets. Thus, at time $k = 0$, the decision variables are

$$\mathbf{U} = \{x_i(k) \text{ for } k = 1, 2,.., K, \text{ and for all assets } i \in I_k\}. \quad (1)$$

Given $x_i(k)$, asset i can reach a set of cells $R_i(x_i(k)) \subset G$ in time period $k+1$ depending on meteorological and oceanographic (METOC) effects at time k and its speed. Thus, $R_i(x_i(k))$ can be a function of k, but does not show its explicit dependence on k for simplicity of notation. The objective function used to select $\mathbf{U}$, via an optimization scheme at $k=0$, is given in (2). Then, on the next day, we solve the problem once again. The optimization algorithm follows a regenerative optimization scheme, i.e., it is in the class of open-loop feedback optimal policies [17]. The optimal policy $\mathbf{U}^*$ is computed over the planning horizon [1, K]; however, of the decisions that are made today ($k=0$), *only* the commands $\{x_i(1)\}$ are to be implemented at $k=1$. Thus, k is only a relative time index, not an absolute one. It should be possible to use the previous optimization results as initial conditions for the
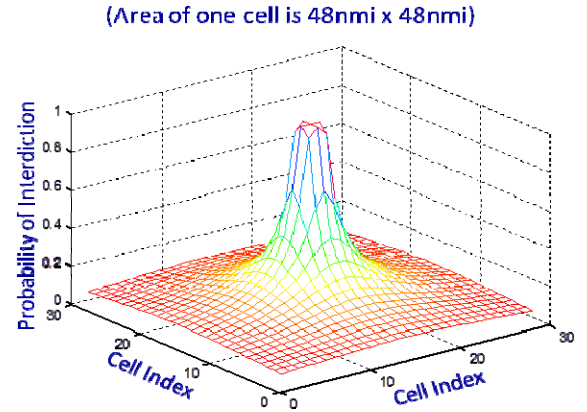

Figure 2: Interdiction probability for an asset located at cell (15, 15)

next period's optimization. Note that our formulation allows a cell to be covered by multiple interdiction assets.

$$\max_{\mathbf{U}} \sum_{k=1}^{K} \left(\gamma^{(k-1)} \sum_{g \in G} PA(g,k) \left[1 - \prod_{i \in I_k} \left(1 - PI_i(x_i(k), g) \right) \right] \right) \quad (2)$$

s.t. $x_i(k+1) \in R_i\left[x_i(k)\right]; x_i(0)$ is given; $i \in I_k; k = 0,.., K-1$

where γ (≤ 1) is the discount factor. The interdiction probability $PI(x_i(k), g)$ is calculated based on the centered 1-D scheme proposed in [3], which takes into consideration the vessel speed, the helicopter speed (if any), and the delay time to launch the helicopter. Following [3], the probability of interdicting a piracy event in cell g within a specified interdiction time τ (typically 30 minutes) is given by

$$PI_i(x_i(k), g) = \begin{cases} \dfrac{r(i, \tau)}{dist(x_i(k), g)}, & r(i, \tau) < dist(x_i(k), g) \\ 1, & r(i, \tau) \geq dist(x_i(k), g) \end{cases} \quad (3)$$

where $dist(x_i(k), g)$ is the Euclidean distance from cell g, where a piracy event takes place, to the asset i's location $x_i(k)$, and $r(i, \tau)$ is the distance that will be covered in a time τ by asset i. It is defined as

$$r(i, \tau) = \begin{cases} v_i \tau & \tau \leq t_i^h \\ v_i t_i^h + v_i^h (\tau - t_i^h) & \tau > t_i^h \end{cases}, \quad (4)$$

where v_i is the speed of asset i, v_i^h is the speed of the helicopter operated by asset i, and t_i^h is the time to launch the helicopter (typically 10 minutes). The interdiction probability using (3) and (4) is illustrated in Fig. 2. Note that the reachable cells $R_i(x_i(k))$ are determined by subtracting the interdiction time τ from the time interval, Δ.

B. Computational Complexity

Consider a map containing a total of $M \times N$ cells. Then, the interdiction problem can be converted into a network model as
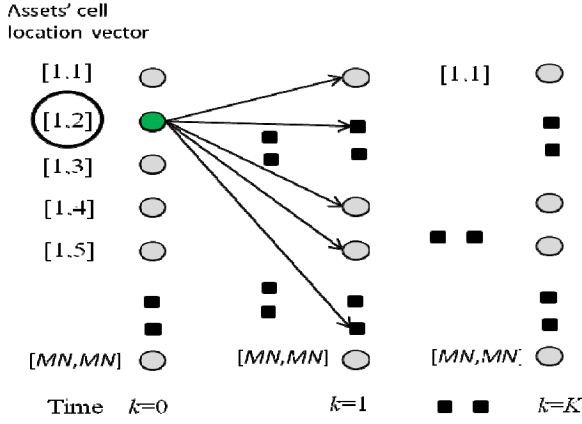
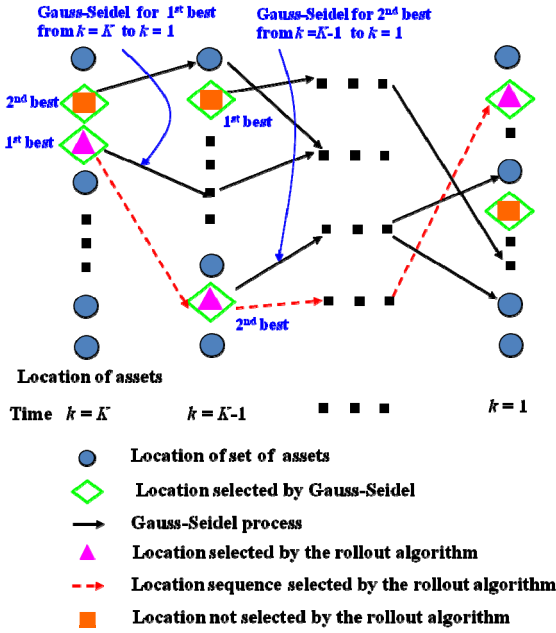Figure 3: Interdiction problem viewed as a very large network

Figure 4: Illustration of rollout coupled with Gauss-Seidel algorithm

shown in Fig. 3, where asset locations correspond to a node, and the directed arcs correspond to the reward (objective function defined in (2)) obtained by moving the assets to their next locations. Here, the optimal solution to the problem in (2) can be obtained using a longest path algorithm. However, as the number of assets increases, the number of nodes in the network increases exponentially as $O(K(MN)^I)$, and hence an *optimal* solution becomes intractable.

C. Gauss-Seidel (*Method of Successive Displacements*) and Rollout Algorithm

The Gauss-Seidel method [15] is an iterative method used to solve a linear system of equations. Since a variable updated in a particular iteration depends upon all previously computed variables and the ordering, the Gauss-Seidel approach is also called the method of successive displacements. Here, we first choose an asset $i \in I_k$ and assuming that the positions of the

remaining assets are fixed, find its reachable cells during a single time period. Then, we allocate the asset to the cell that provides the best reward and repeat the above process over the planning horizon. This process of selecting an asset and finding its positioning over the planning horizon is repeated until the sum of interdiction probabilities given in (2) converges to a (typically local) optimum.

The Gauss-Seidel concept is easily applied to the interdiction problem. In contrast to the approach using a network model, the Gauss-Seidel method requires one to explore a number of nodes that is linear in the number of cells, the number of assets, and the time horizon $O(KIMN)$ per iteration. The algorithm makes use of the maximum reachable range, $R_i(x_i(k))$ associated with each asset during each time period to further reduce the number of nodes to be explored. The maximum reachable range, $R_i(x_i(k))$, is a function of the current location $x_i(k)$ of asset i, meteorological and oceanographic (METOC) effects at time k, and asset speed.

Another advantage of our approach is that the rollout algorithm can be applied, since the computations in the proposed approach are serial. Rollout algorithms are a class of suboptimal methods, which are capable of improving the effectiveness of any given heuristic through sequential application. This is due to the fact that rollout can be viewed as a single step of the classical policy iteration method used to solve Markov decision problems [16][17], wherein one starts from a given easily implementable and computationally tractable policy, and then tries to improve on that policy using online learning and simulation. The attractive aspects of rollout algorithms are its simplicity, broad applicability, and suitability for online implementation. The rollout algorithm, combined with the Gauss-Seidel approach, is shown schematically in Fig. 4.

III. SURVEILLANCE PROBLEM

A. Surveillance Problem

In general, P-3s are used for large ocean surveillance and generally perform one surface search mission per day. The assets are assigned to predefined search regions in a "box," where the actual search time allowed for the asset is determined by excluding the transit time from the time interval, Δ. In this section, we formulate the surveillance problem, where the area not covered by interdiction assets is partitioned into search regions having rectangular shapes. A set S_k of available surveillance assets at time k is assigned to the partitioned regions to maximize the discounted cumulative sum of detection probability over the planning horizon, $k=1,2,..,K$.

A search region assigned to a surveillance asset s at time k is given by the set of cells $A_s(k)$, which is a rectangular subset of cells in G. It is defined by two coordinates comprising a longitude and latitude, $(long_1(s), lat_1(s))$ and $(long_2(s), lat_2(s))$. Here, $(long_1(s), lat_1(s))$ are the coordinates of the lower left cell in the search region $A_s(k)$, while coordinates $(long_2(s), lat_2(s))$ refer to the location of the upper right cell in $A_s(k)$.

Each surveillance asset, $s \in S_k$, can have different capabilities measured in terms of the sweep width $w_s(k)$ and the speed $v_s(k)$. Note that the effective sweep width $w_s(k)$ of asset s is a function of METOC conditions in the region at a particular time k. Let the probability map of pirate presence be
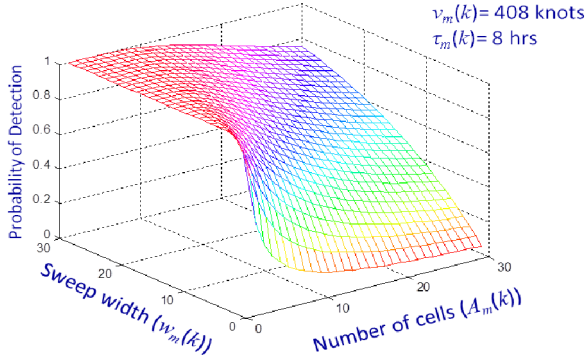
$v_m(k)$= 408 knots
$\tau_m(k)$= 8 hrs

Figure 5: Probability of detection for surveillance asset vs. sweep width and number of cells

Expandable cell   Expanded cell
Initially selected cell   Selected cell

Figure 6: Method of cell growth for maintaining rectangular shape

Table I: Vessel characteristics

Asset ID	Max Speed (km/h)	Range (km)	Home Base	Asset Types Carried	# of Assets Carried	Surface Radar Range (km)
CG-55	60	6100	US	H	2	92
F-208	56	4000	DE	H	2	46
DD-101	56	4000	JP	H	1	85
F-111	50	5600	NZ	H	1	33
DDG-175	56	4200	JP	-	-	125
F-570	61	6000	IT	H	2	150
LPD-17	41	-	US	B	4/2	410
DDG-72	56	4100	US	H	1	104
F-79	52	8900	UK	H	1	200

Table II: Aircraft characteristics

Asset ID	Max Speed (km/h)	Range (km)	HomeBase	Sweep Width	Launch Time Delay (min)
P-3	750	4400	US	2	-
V-22A	463	1627	US	2	-
SH-60B	270	834	US	1	10
MH-60R	270	834	US	1	10
SH-60J(K)	270	834	JP	1	10
MK-88A	324	528	DE	1	10
SH-2G	261	1080	NZ	1	10
AB-212	220	460	IT	1	10
CH-46	265	296	US	1	10
MH-1	300	1389	UK	1	10
UAV-SE	139	100	US	2	-

such that $PP(g, k)$ denotes the probability that at least one pirate is in cell g at time k. Let the amount of time asset s, $s \in S_k$, spends in the assigned search region A_s during time step k be given by $\tau_s(k)$. Following [22][23], the probability of detection of asset s assigned to a set of cells $A_s(k)$ is given by

$$PD_s\big(A_s(k),k\big)$$
$$= \sum_{g\in A_s(k)} PP(g,k)\times\left(1-\exp\left(\frac{-w_s(k)*v_s(k)*\tau_s(k)}{a_c\,|A_s(k)|}\right)\right), \quad (5)$$

where a_c is the area of a cell and $|A_s(k)|$ is the number of cells in the search region, $A_s(k) \subset G$. Fig. 5 shows how the probability of detection changes as the sweep width and number of cells are varied. Now, the surveillance problem can be succinctly written as follows.

$$\max_{\{A_s(k), s\in S_k\}} \sum_{k=1}^{K} \gamma^{(k-1)} \sum_{s\in S_k} PD_s(A_s(k),k)$$
$$\text{s.t. } A_i(k)\cap A_j(k)=\varnothing, \quad (i\neq j)\in S_k \quad (6)$$
$$A_i(k) \text{ has rectangular shape, } \forall i\in S_k$$

However, the partitioning problem is computationally intractable. This motivates us to develop a heuristic algorithm

to "grow" search areas. The process for growing rectangular search areas is shown in Fig. 6. The key idea of the algorithm is to grow the current search region by selecting its expandable cells and reshape it to a new rectangle. The algorithm initially starts with a single randomly selected cell, along with potential expandable cells (denoted by circles in Fig. 6). Then it randomly selects one of these expandable cells, updates the shape while preserving the rectangularity (as indicated by an arrow in Fig. 6), and then revises the expandable cells for the new rectangle. This process is continued until there are no more expandable cells.

Using the proposed algorithm, we first generate $N \geq |S_k|$ search regions, where at most one (including zero) sensing asset can be assigned to a search region and only one search region can be assigned to an asset. The above problem of allocating $|S_k|$ sensing assets among N search regions is a one-to-one asymmetric assignment problem. The auction algorithm [21] is one of the most efficient methods for solving the one-to-one assignment problems. It consists of a bidding phase and an assignment phase, where an optimal assignment is found by employing a coordinate descent method on the dual function. In order to use auction, one needs to consider the $|S_k| \times N$ matrix of detection probabilities for each asset-region pair as in (5). The process of creating rectangles and solving the concomitant assignment problems continues until the discounted cumulative detection probability over the planning horizon converges. A discussion on choosing the appropriate number of the search regions is presented in the next section.

IV. COMPUTATIONAL RESULTS

A. Scenario and Results

Here, the AOR is partitioned into cells corresponding to the available METOC forecasts, 0.8° square, or roughly 48 nmi ×
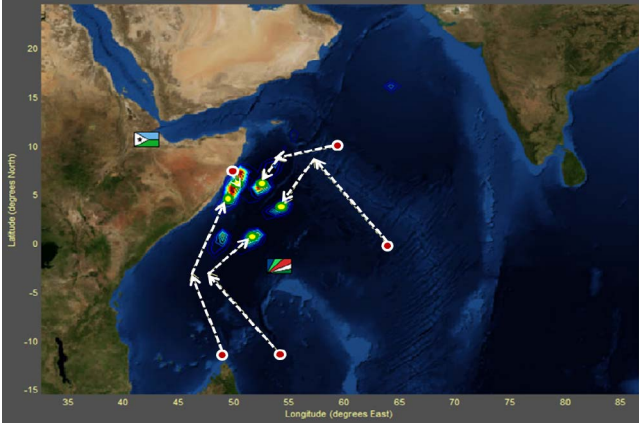
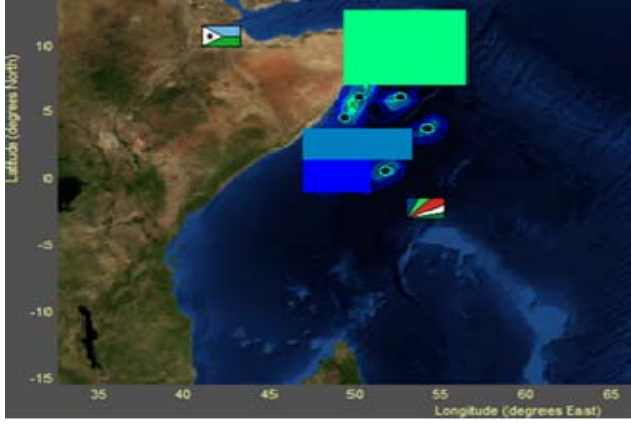Figure 7: Probability map and paths of interdiction vessels

Figure 8: Partitioned uncovered area using three surveillance assets (P-3) at k=2

48 nmi, and indexed by $g \in G$. Specific locations for surface-based assets (e.g., P-3 aircraft) are used in conjunction with the 43×51 cells of AOR. There are two types of assets (see Table I and Table II):

- Vessels (frigates and destroyers): The vessels can carry out interdiction only or interdiction and surveillance missions.

- Land-based aircraft (P-3s operating out of two distinct bases in the AOR): These aircraft only have surveillance capability.

In this example five equally capable interdiction vessels and three surveillance assets are assigned to conduct a counter-piracy mission. We first solved the interdiction problem, i.e., we determined the locations of interdiction assets only, over the planning horizon K=2 ($k\varDelta$=12, 24 hours). Typically, K=6, but we have chosen K=2 for clarity in displaying the trajectories of the interdiction assets. The interdiction probability $PI(x_i(k), g)$ is calculated using (5) based on the scheme proposed in [3]. Fig. 7 shows the paths of the interdiction vessels over the PARS maps, where the initial and final locations are indicated by the red and yellow cells, respectively. The path for each interdiction vessel is indicated by arrows. The results show that all vessels move to the cells with high probability of pirate attack as time progresses.
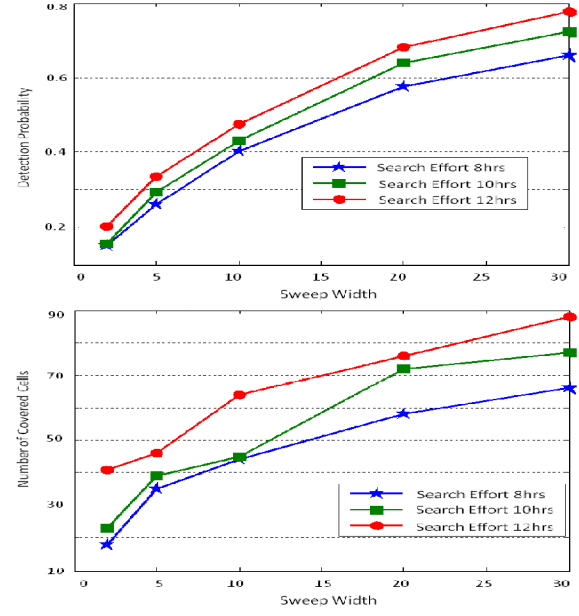

Figure 9: Detection probability and number of covered cells vs. sweep width
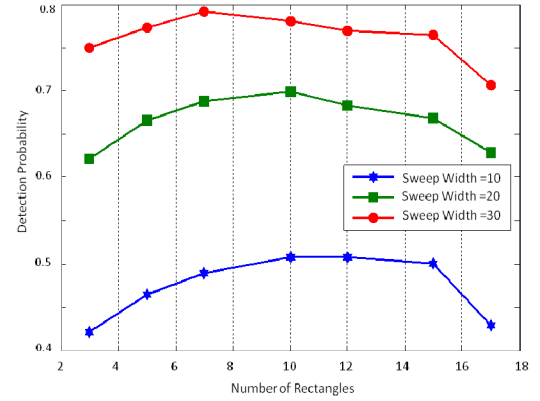

Figure 10: Detection probability vs. number of rectangular search regions

Then, we solved the surveillance problem with three land-based surveillance aircrafts (e.g., P-3's), where we needed to partition uncovered areas into search regions (rectangular regions) at a given time epoch. Given the partitioned regions, the detection probabilities are computed by using (5). Fig. 8 shows the search boxes at time k =2 using the proposed approach. The results show that the surveillance assets are assigned to cover regions of high pirate presence probability.

B. Key Research Questions

1) Simulating Surveillance Scenario with different sweep widths: Fig. 9 shows the detection probability and the number of cells covered as a function of the sweep width for different search efforts. As the sweep width and search time decrease, it is better for the assets to focus on small search areas with high probability of pirate presence. On the other hand, as the sweep width and search time increase, it is better for the assets to cover a larger area as well as the high probability areas, which is commensurate with our intuition. Fig. 10 shows the detection probability as a function of the number of rectangles

Table III: Number of detection number for three surveillance assets
(P-3's) with 10nm sweep width

	Number of targets = 1		Number of targets = 2	
	Update 1	Update 2	Update 1	Update 2
East - West	1158 / 1491	1145 / 1489	1852 / 2932	1871 / 2957
North - South	1187 / 1570	1143 / 1504	1893 / 3071	1943 / 3170
IG	789 / 2856	808 / 2877	1408 / 5591	1341 / 5571

for different sweep widths. It shows that there is an optimal number of rectangles for each sweep width, although it is nearly flat at high sweep widths. Thus, it is important to choose the appropriate number of search regions in our algorithm, especially during low visibility conditions.

2) Surveillance with different search strategies: Once the surveillance assets are assigned to the partitioned search areas at the operational level of planning (e.g., in a maritime operations center), the dynamic search paths at the tactical level are determined. Here, we explore three such search strategies. First, we consider the East-West (E-W) search strategy, where we begin the search from the lower left-most cell and move horizontally to the right-most cell. If there is no cell with non-zero probability of target presence, then it moves to the upper right-most cell and continues horizontally towards the left-most cell. Similarly, the asset follows the North-South (N-S) path in a N-S search strategy. The information gain (IG) heuristic selects the best asset allocation at each time epoch that maximizes the sum of information gains over all the cells [5][16]. In evaluating the search strategies, we assume that the velocity of each target (skiff) is 30 nmi/hr and it can randomly move 8 cells in 12 hours, that is, transitions are equally likely among the neighboring cells (including staying in the same cell) in each subinterval of duration approximately 1.5 hours. Here, we consider two scenarios with a single target and two targets within a rectangular search region. For each of the above scenarios, two update rules are considered; in update 1, the probability map is updated only for the visited cells; in update 2, we update non-visited cells as well. In both caess, we normalize the pirate presence map to equal the expected number of targets. Based on 3000 Monte-Carlo runs, Table III shows the performance of the three search strategies mentioned above. Let us denote by D_j the total number of detections over the j^{th} run. We define two parameters A and B based on $\{D_j\}$. The parameter A is the cumulative number of detections over all runs, while the parameter B is the number of runs with at least one detection event. Formally,

$$A = \sum_j D_j,$$
$$B = \sum_j I(D_j),$$
(7)

where $I(D_j)$=1, if $D_j > 0$, otherwise $I(D_j)$=0. The values in Table III are presented in the form of B/A, the ratio of number of runs with at least one detection event and the cumulative number of detections over all runs. The smaller this ratio is for a search strategy, the greater is the persistence in tracking the target. The result shows that IG-based search strategy is very effective on this metric.
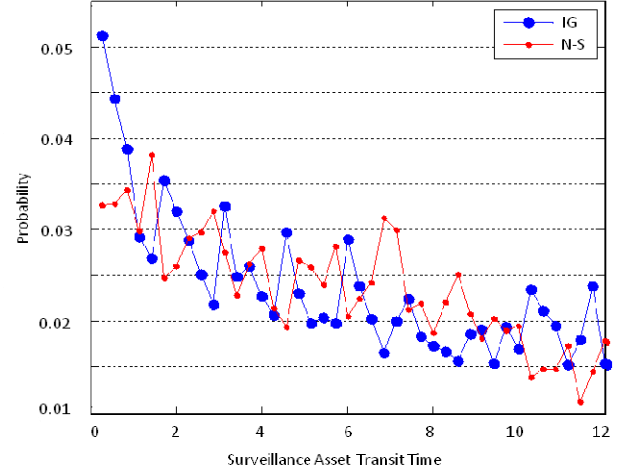

Figure 11: Distribution of first detection over time interval (Δ=12hrs)

3) Mean-time-to-first detection: To examine this idea, we consider a scenario with a single target and three surveillance assets, each with a sweep width of 20 nmi. In Fig. 11, we plot the distribution of time-to-first-detection over the time interval (Δ = 12 hrs) for the N-S and IG-based search strategies over 3000 Monte Carlo runs. The E-W search startegy is not included here as it has similar characteristics to the N-S search strategy. The time-to-first-detection appears to follow an exponential distribution. In particular, the IG-based strategy has smaller time-to-first-detection than the non-adaptive N-S search strategy. That is, the pirate presence map, coupled with an IG-based search strategy, makes the search (and the concomitant interdiction) process more effective by decreasing the mean-time-to-detect an asymmetric threat .

V. CONCLUSIONS

This paper developed a set of optimization algorithms for allocating interdiction and surveillance assets within a counter-piracy mission environment. In order to overcome the curse of dimensionality of dynamic programming recursion, we proposed a method of successive displacements and rollout concepts for solving the interdiction problem. For the surveillance problem we proposed a partitioning algorithm, where each search region is grown independently subject to the region's shape constraints. The performance of tactical search strategies (North-South, East-West, and information gain heuristic) was compared using the cumulative detection probability as a performance metric.

The analytical models developed in this paper provide a systematic framework to take PARS-like information, such as probabilities of pirate presence and risk of attack, and use it effectively in the subsequent process wherein assets are to be allocated and positioned over time to best thwart potential attacks. As a result of this modeling work operational mission planners will have the ability to optimize how they allocate their limited counter piracy assets over a large geographical area. We have developed and demonstrated models for optimizing within interdiction and surveillance missions, and have a basis to adjust the asset allocations to maximize for

either of these objectives at any given time in the face of prevalent weather conditions and sea states.

Future extensions to our model will incorporate realistic counter-piracy mission environmental features (e.g. ensemble forecast uncertainties associated with PARS), dynamic asset status, as well as observed detection and interdiction events. Also, there are a number of other search strategies for possible consideration. These includes locating surveillance assets to: 1) Perform alert/confirm type surveillance, where an alert triggers a confirmation cycle in which multiple measurements are taken in the alert cells [5]; 2) Sample cells with the highest probability of pirate presence (multi-armed bandit index rule); and 3) Minimize the sum of the variances of the pirate presence probabilities over all cells. There are also other performance metrics, such as probability gain, impact, Bayesian diagnosticity and log diagnosticity [6] that are worth exploring. Additionally, we plan to extend our formulation to consider the effects of uncertainty in probability maps and weather impact on the dynamics of asset motion. In particular, we plan to explore approximate dynamic programming (ADP) techniques for overcoming the curses of dimensionality, viz., the state space explosion, randomness in probability maps and weather-impacted asset motion, as well as the large decision space associated with locating interdiction and surveillance assets in a large AOR [17, 31].

REFERENCES

[1] MSNBC.com, "Anti-pirate software tracks risks on high seas: U.S. Navy computer model tries to predict where criminals will strike next," May 23, 2011.

[2] ICC Commercial Crime Service, "Attacks off the coast drive piracy to record high, reports IMB", April, 2012.

[3] R. Mirshak, "Ship Response Capability Models for Counter-Piracy Patrols in the Gulf of Aden," *DRDC CORA TM* 2011-139, September 2011.

[4] G. K. Jabeur, M. Allouche, H. Wehn, J. Happe, "Application of Search Theory for Large Volume Surveillance Planning", pp 1-8, 11th *International Conference on Information Fusion*, 2008.

[5] A. Hero, D. Castanon, D. Cochran and K. Kastella (Eds.), "Foundations and Applications of Sensor Management", *Springer*, 2008.

[6] J. D. Nelson, "Towards a Rational Theory of Human Information Acquisition," Chapter 9 in N. Chater and M. Oaksford (Eds.), *Probabilistic Mind*: *Prospects for Bayesian Cognitive Science*, Oxford, 2008.

[7] G. H. Martins, "A new branch-and-bound procedure for computing optimal search paths", *Master's Thesis*, Naval postgraduate school, 1993.

[8] H. Lau, "Optimal Search in Structured Environments," *PhD Thesis*, University of Technology, Sydney, 2007.

[9] W. An, M. Mishra, C. Park, and K. R. Pattipati, "An integrated asset allocation and path planning method to search for targets in a dynamic environment", 15th International Command and Control Research and Technology Symposium, Santa Monica, CA, June 2010.

[10] J. O. Royset, H. Sato, "Route Optimization for Multiple Searchers", Naval *Research Logistics*, Vol. 57, 2010

[11] S. Nair, K. R. Chevva, H. Owhadi, and J. Marsden, "Multiple Target Detection using Bayesian Learning," *Proceedings of the 48th IEEE Conference on Decision and Control*, pp. 8549-8554 , Shanghai, P.R. China, 2009

[12] L. A. Slootmaker, "Countering piracy with the next-generation piracy performance surface model", *Master's Thesis*, Naval postgraduate school, 2011.

[13] L. Esher, S. Hall, E. Regnier, P. J. Sanchez, J. A. Hansen, and D. Singham "Simulating pirate behavior to exploit environmental information" *Proceedings of Winter Simulation Conference,* pp. 1330-1335, 2010.

[14] K. Byers, "Situational Awareness for surveillance and interdiction operations (SASIO): Tactical installation protection", *Master's Thesis*, Naval postgraduate school, 2010.

[15] D. P. Bertsekas, "Nonlinear Programming", *Athena Scientific*, Belmount, MA, 2003.

[16] W. An, S. Singh, K. R. Pattipati, S. S. Gokhale, and D. L. Kleinman, "Dynamic scheduling of multiple hidden Markov model-based sensors", *Journal of advances in Information fusion*, vol. 3, pp. 33-49, 2008.

[17] D. P. Bertsekas, "Dynamic Programming and Optimal Control/ Approximate Dynamic Programming*", Athena Scientific,* 2007.

[18] J. Eagle,"The optimal search for a moving target when the search path is constrained", *Operations Research*, 32, 1107-1115, 1984.

[19] J. F. Caldwell, "Investigation and implementation of an algorithm for computing optimal search paths", *Master's Thesis*, Naval postgraduate school, 1987.

[20] A. G. Santos, "Using multiple searchers to locate a randomly moving target" *Master's Thesis*, Naval postgraduate school,1993.

[21] D. P. Bertsekas, "Network Optimization: Continuous and Discrete Models", *Athena Scientific*, Belmount, MA, 1998.

[22] B. O. Koopman, "Search and Screening: General Principles with Historical Applications", *Pergamon Press*, New York, 1980.

[23] U.S. Coast Guard Research and Development, "Review of search theory: advances and applications to search and rescue decision support", *Final Report*, 2001.

[24] W. An, M. Mishra and K.R. Pattipati, "Impact of Oceanographic Environment on ASW Mission Planning and Risk Assessment," Technical Report, ECE-Cyb-2012-01, January 2012.

[25] N. Rozen, "Sensor-Interceptor operational policy optimization for maritime interdiction missions", Master's Thesis, Naval postgraduate school, 2009.

[26] C. D. Marsh, "Counter Piracy: A repeated game with asymmetric information", Master's Thesis, Naval postgraduate school, 2009

[27] M. Kress, J.O. Royset, N. Rozen, "The Eye and the Fist: Optimizing Search and Interdiction", *European Journal of Operational Research*, 2012 (accepted).

[28] One Earth Future Foundation, "The economic cost of Somali Piracy", *Working Paper*, 2011.

[29] Insurance Journal News "2011 Piracy Attacks Totaled 439; 275 off Somalia: ICC/IMB Report", January, 19, 2012.

[30] ASI Global, "Maritime Response", Piracy report, 2010.

[31] W. B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Wiley-InterScience, Hoboken,NJ, 2007.

[32] C. D. Marsh, "Estimations of atmospheric conditions for input to the radar performance surface", *Master's Thesis*, Naval postgraduate school, pp 1-2, 2007